



**HAL**  
open science

## Continuous Authentication Leveraging Matrix Profile

Luis Ibanez-Lissen, Jose Maria de Fuentes, Lorena Gonzales-Manzano, Nicolas Anciaux

► **To cite this version:**

Luis Ibanez-Lissen, Jose Maria de Fuentes, Lorena Gonzales-Manzano, Nicolas Anciaux. Continuous Authentication Leveraging Matrix Profile. ARES 2024 - The 19th International Conference on Availability, Reliability and Security, Jul 2024, Vienne, Austria. hal-04663471

**HAL Id: hal-04663471**

<https://inria.hal.science/hal-04663471v1>

Submitted on 28 Jul 2024

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

# Continuous Authentication Leveraging Matrix Profile

Luis Ibanez-Lissen

luibanez@pa.uc3m.es

Universidad Carlos III de Madrid

Leganes, Madrid, Spain

Lorena Gonzalez-Manzano

lgmanzan@inf.uc3m.es

Universidad Carlos III de Madrid

Leganes, Madrid, Spain

Institut Polytechnique de Paris

Palaiseau, France

Jose Maria de Fuentes

jfuentes@inf.uc3m.es

Universidad Carlos III de Madrid

Leganes, Madrid, Spain

Inria

Palaiseau, France

Nicolas Anciaux

nicolas.anciaux@inria.fr

Inria, UVSQ, U. Paris-Saclay

Palaiseau, France

## ABSTRACT

Continuous Authentication (CA) mechanisms involve managing sensitive data from users which may change over time. Both requirements (privacy and adapting to new users) lead to a tension in the amount and granularity of the data at stake. However, no previous work has addressed them together. This paper proposes a CA approach that leverages incremental Matrix Profile (MP) and Deep Learning using accelerometer data. Results show that MP is effective for CA purposes, leading to 99% of accuracy when a single user is authorized. Besides, the model can on-the-fly increase the set of authorized users up to 10 while offering similar accuracy rates. The amount of input data is also characterized – the last 15 s. of data in the user device require 0.4 MB of storage and lead to a CA accuracy of 97% even with 10 authorized users.

## CCS CONCEPTS

• Security and privacy → Authentication; Privacy protections.

## KEYWORDS

Continuous authentication, privacy, data minimization, storage limitation, matrix profile, deep learning

## ACM Reference Format:

Luis Ibanez-Lissen, Jose Maria de Fuentes, Lorena Gonzalez-Manzano, and Nicolas Anciaux. 2024. Continuous Authentication Leveraging Matrix Profile. In *Ares '24: Availability, Reliability and Security Conference, July 30–August 02, 2024, Vienna, Austria*. ACM, New York, NY, USA, 13 pages. <https://doi.org/XXXXXXX.XXXXXXX>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

*Ares '24, July 30–August 02, 2024, Vienna, Austria*

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-XXXX-X/18/06

<https://doi.org/XXXXXXX.XXXXXXX>

## 1 INTRODUCTION

The amount of information being shared on the internet <sup>1</sup> is promoting new ways of data mining and novel Artificial intelligence that can be used to enhance the security of users and information. The proliferation of IoT devices [4] has given to the academia the access to high quality data and sensors able to capture environmental conditions and users behaviors.

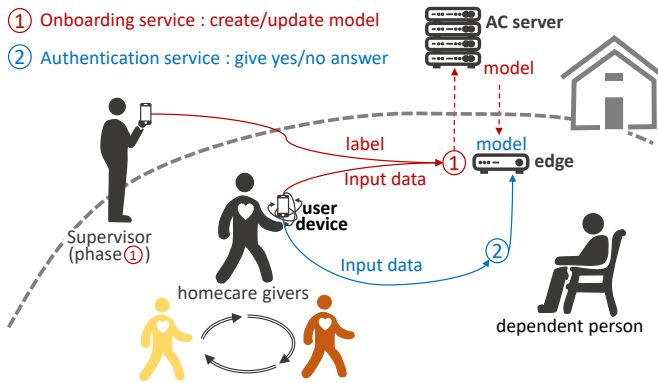
By analyzing this data, it is possible to identify unique patterns that Continuous Authentication (CA) systems can exploit for recognizing individuals. These patterns can be processed in real-time to ensure the authenticity and security of users' interactions with machine learning algorithms.

Current existing (and commercial) applications of CA are based on, for instance, motion sensors which provide information about a user's physical activities [2] and keystroke dynamics to recognize individuals based on their typing rhythm and pressure [37]. Additionally, biometric data, such as EEG or ECG readings, provide additional features for identifying unique patterns related to brain activity or heart rate variability [39, 44].

**Commercial use and advanced scenarios:** In real-world scenarios, CA systems are increasingly integrated into ubiquitous environments like smart-homes, where IoT devices authenticate users seamlessly [34], exemplifying the potential in sectors such as healthcare. For example, in a scenario (see Figure 1) where healthcare services that are provided to a dependent individual by one or multiple assistants from an outsourced company, the CA system, controlled by the dependent individual or referent professionals, must authenticate each assistant.

Two primary challenges arise in this scenario. Firstly, *Privacy compliance* is crucial, given the handling of highly sensitive data, necessitating adherence to regulations like GDPR or CCP/CPRA. This makes CA approaches relying on sending raw data to third-party servers for processing authentication proofs (e.g., [10]) unsuitable. Secondly, the *Adaptive user* challenge involves accommodating the changing identity of the authorized users (e.g., assistants in the example above), requiring the CA mechanism to be flexible enough in this regard. Note that new authorized users may be added to

<sup>1</sup><https://www.statista.com/statistics/871513/worldwide-data-created/>, last access January 2024



**Figure 1: Continuous authentication in a home care scenario.**

the system over time. Thus, it must enable these changes while remaining in operation.

This scenario is representative of broader applications facing both "privacy" and "user adaptation" issues. Previous works have already explored various privacy-preserving approaches to address the first challenge. Sanchez *et al.* [45] introduced a multi-device CA system where distributed data is collected from various devices, while Wazzeah *et al.* [54] employed federated learning to distribute predictions between a set of local and online models, minimizing the amount of raw data exchanged between devices and servers. Nonetheless, this approach could be vulnerable to model inversion attacks [17], where an attacker may extract sensitive information from the distributed models.

In what comes to the second challenge, Chauhan *et al.* [9] proposed using breathing patterns for CA and explored incremental learning as a way to address scalability and model adaptation to changes in the data. Shen *et al.* [47] with a different approach, makes use of touch dynamics as authentication proofs. These methods focus on scalability and adaptability but lack sufficient data protection measures, limiting their use in realistic settings such as our healthcare scenario.

The inherent conflict in addressing these challenges lies in the contradictory nature of their requirements: *Privacy compliance* necessitates minimal personal data collection and limited storage, whereas the *Adaptive user* challenge often requires extensive personal data collection and diverse data sources.

To address these issues, our paper introduces a novel CA approach that combines Matrix Profile (MP) [64] and Deep Learning (DL) [30]. MP is particularly adept at addressing the *Privacy compliance* challenge as it processes a transformed representation of raw time series data, which aligns better with the GDPR principles of Data minimization and Storage limitation<sup>2</sup>. However, no previous work has applied it in this context. Meanwhile, DL's ability to learn and adapt the behavioral profiles of authorized users make it suitable for addressing the *Adaptive user* challenge. Our

<sup>2</sup>This transformation can be computed locally and helps in reducing the risk of exposing personal information, although formally quantifying the extent of this non-reversibility is left for future work. In particular, MP can hide local patterns of a sensitive time series for anonymization [12], while preserving the utility of the original time series in certain contexts.

**research question** and contribution thus focus on the feasibility of performing efficient continuous authentication for adaptive users in a privacy-compliant manner:

- We explore the use of a MP variant (incremental MP [63]) to identify features for CA for a single user.
- Different buffer size and windows size, have been set to measure the quality of the Deep learning models.
- We measure how well the mechanism scales when adding new authorized users.
- We characterize the amount of data that is to be stored by all the devices at stake.
- The code and models are publicly released to foster further research.

*Paper organization.* Section 2 describes the model and problem statement, illustrated by means of a concrete scenario. Next, Section 3 presents the two building blocks on which our proposal is based, namely MP [64] and DL [30]. Section 4 describes the holistic solution we propose, which is articulated as a combination of these two building blocks. The preparation of the experiments is dealt with in Section 5, while assessment is presented in Section 6. Related work is analysed in Section 7. Finally, the paper is concluded in Section 8 and future research directions are pointed out.

## 2 PROBLEM STATEMENT

We first introduce the general continuous authentication problem in Section 2.1. Then, using a real-world healthcare scenario, we present the main assumptions in Section 2.2. Finally, the problem statement including the precise goals at stake are formulated in Section 2.3.

### 2.1 Continuous Behavioural Authentication

A biometric authentication system (BAS) is a cybersecurity mechanism that allows a user to authenticate by making use of their physiological (iris, hands, fingerprints) or behavioural features [41, 53]. Authentication proofs against BASs can be based on three main principles: "Something you know," "Something you have," "Something you are". Each of them have different pros and cons, however, the third principle have been gaining more and more attraction due to its balance between usability and accuracy [1]. The generalization of IoT devices and smartphones have opened the opportunity to explore sensor based applications where the BASs can learn to recognize unique features and patterns that allows to recognize individuals by the way they interact with their environment and devices [8].

Continuous biometric authentication faces the challenge of distinguishing a legitimate user from adversaries or illegitimate users.

In this context, the authentication system  $\mathcal{S}$  must learn to recognize and understand the unique authentication proofs,  $p$ , prepared by any of the authorized users  $\mathcal{U} = u_1, u_2, \dots, u_N$ . Potential adversaries represent individuals who may attempt unauthorized accesses. Thus, the set of adversaries is  $\mathcal{A} = \mathcal{W} \setminus \mathcal{U}$  where  $\mathcal{W}$  is the universe of human subjects.

The verification function  $V(p)$  is responsible for assigning likelihood scores to each authentication proof  $p$ . These scores determine whether the proof belongs to a user  $u_i \in \mathcal{U}$  or an adversary within  $\mathcal{A}$ .

## 2.2 Scenarios and Assumptions

The representative scenario depicted in the introduction (see Figure 1) is inspired by a real-world application of medical and social care monitoring practiced in France [6, 27]<sup>3</sup>. In this particular scenario, dependent individuals receive at-home visits from professionals who provide medical and social care support. These professionals use a smartphone application (referred to as the "user device" in the figure) to connect to a homebox (denoted as "edge"), allowing them to access (both read and write) the patient's medical and social data. It is important to note that access rights are differentiated based on the role of the caregiver.

Continuous authentication of caregivers is crucial, for instance, to enable audits that inform family members which professional visited the elderly person at home and when. The initial visit to a patient's home by a new caregiver requires a supervisor's presence to introduce the professional to the dependent individual. During this visit, the supervisor can label data produced by the caregiver's user device to facilitate the caregiver's onboarding and, if necessary, adjust the authentication model using a remote server designed for training (phase ① in the figure). For subsequent visits (phase ②), the caregiver can be recognized through the application (user device) and the homebox (edge) with local processing.

There are then two authentication models that can be considered and that make sense in such scenarios:

- **Single authorized user  $u_1$ .** In this model, a single user is authenticated. This is the case of a handheld device being used to gather data to authenticate a user in a given sensitive space (e.g., a designated homecare giver in our context, the access to a corporate data processing center, etc.). Thus,  $\mathcal{S}$  follows a one-vs-rest ( $u_1$  vs.  $\mathcal{A}$ ) fashion, distinguishing one positive class from the rest, creating a binary classification between legitimate and non-legitimate users.
- **Increase authorized users  $u_1, \dots, u_n$ .** In this case, the new authorized users are added over time. Thus, after training with data from  $u_1$ , further authorized users  $u_{i \neq 1}$  are iteratively added to  $U$  on request (e.g., as a result of a hiring process). Here,  $\mathcal{S}$  has to consider several positive classes (to distinguish between  $u_i$ ) as well as a single null negative class (for  $\mathcal{A}$ ).

It must be noted that an intermediate scenario (i.e., a single user whose identity changes from time to time) could be possible. However, we focus on these two scenarios as they illustrate the effect of adding new identities to the system.

In this regard, three main assumptions will be undertaken:

- First,  $\mathcal{S}$  relies on some knowledge to characterize the behavior of  $\mathcal{A}$ . This is reasonable since myriads of behavioral datasets are publicly available.
- Second, any  $u_i$  which is to be considered as legitimate by  $\mathcal{S}$  counts on a certain period of trusted usage. In this way,  $\mathcal{S}$  is able to properly extract the behavioral patterns of that user. In the real world, this may happen during the employee onboarding process, as described before.

- Lastly,  $\mathcal{S}$  is not trusted by any  $u_i \in \mathcal{U}$ . Therefore, it is assumed that  $\mathcal{S}$  may leak the behavioral patterns acquired from  $u_i$  to any (potentially malicious) third party.

## 2.3 Overall Objectives

The aim of this proposal is the design and implementation of Continuous Authentication in environments where user dynamics and trust levels may vary as exemplified by the homecare scenario. This scenario, where several caregivers access the medical and social data of a dependent person, underlines the importance of CA to manage data access in a secure and efficient way. This may be translated into a set of concrete objectives:

- O1 Privacy compliance (data minimization, storage limitation).** Compliance with privacy regulations such as GDPR or CCPA/CPRA is paramount. Specifically, the *Data Minimization*<sup>4</sup> principle ensures that only necessary data is collected. At the same time, retention policies require that personal data be kept no longer than necessary<sup>5</sup>. This is crucial in our scenario, where caregivers' interactions with the patient's data must respect these privacy mandates.
- O2 Accuracy.** The system  $\mathcal{S}$  must maintain a high level of accuracy, minimizing both false rejection and false authentications. In the context of our scenario, this means ensuring that caregivers are accurately authenticated, allowing seamless authentication, audit and/or secure access to patient data without compromising security or convenience.
- O3 User scalability.** The system  $\mathcal{S}$  should accommodate the addition of new users  $u_{i \neq 1}$  over time, reflecting the dynamic nature of care teams. This scalability is essential to adapt to changes in care providers without compromising the security or efficiency of the authentication process.

The problem we address in this paper is therefore to propose a solution for continuous authentication that meets these objectives –privacy compliance, accuracy and user scalability– in the specific context of home care scenarios. Given the inherent conflicts between these objectives, existing authentication solutions fall short and require a novel approach as discussed in section 7.

## 3 BUILDING BLOCKS

This section introduces the main notions related to the proposal to meet both privacy and performance goals, namely the notions of Matrix profile (Section 3.1) and two notions of Artificial Intelligence –Deep learning and lifelong learning (sections 3.2 and 3.3, respectively). These two building blocks will then be combined in the next section into an end-to-end solution.

### 3.1 Matrix profile

Matrix profile (MP) is an efficient and scalable algorithm for time series subsequence all-pairs-similarity-search [56], providing an exact solution to the discovery of repeated patterns, *motifs*, and

<sup>4</sup>Data Minimization is defined in GDPR Article 5.1.c as the fact that data must be "adequate, relevant and limited to what is necessary in relation to the purposes for which they are processed ('data minimisation')"

<sup>5</sup>According to GDPR Article 5.1.e, personal data shall be "kept in a form which permits identification of data subjects for no longer than is necessary (...) to safeguard the rights and freedoms of the data subject ('storage limitation')".

<sup>3</sup>See the first cited reference section 4.4 paragraph entitled "A concrete ES-PDMS instance", and the second one section 4.

novel or different ones, *discords*; it has shown to be a useful tool for data mining techniques [60].

When computed for a given time series  $T = t_0, \dots, t_n$ , MP returns a distance profile consisting of a vector of distances between all pairs of subsequences  $(t_i, \dots, t_{i+m-1})$  of a size  $m$ , referred to as window size (WS). Distance between subsequences can be calculated with any distance function. However, a standard function is the z-normalized Euclidean distance [11, 35] since it shows better results when used to compare time-series [23].

Matrix profiles can be computed in two settings: offline and online. In an offline setting, MP is calculated with a fixed amount of data when the streaming has stopped; thus, being non-incremental requires re-passing the existing dataset and recomputing the whole MP distances matrix. On the contrary, in online settings, the distance between motifs and discords, as well as the distance matrix, is calculated in an incremental fashion. Thus, it is typically referred to as incremental MP [61, 62, 64].

### 3.2 Deep convolutional networks

Deep convolutional networks (DCN) are a particular type of neural networks [30]. They are formed by a number of layers and convolutional operations, each one having a set of neurons. They are considered to be within the Deep Learning techniques when the deepness of the neural network is more than three, counting input and output layers [29]. DCNs have already been successfully applied in a wide range of image recognition and classification tasks [26, 48, 51, 58].

The core operation of DCNs is the so-called convolution, a linear operation used to extract relevant features from a grid-like input data of any dimension, such as time-series (1-D), Images (2-D), and Objects (3-D).

For 3D-like inputs, architectures and operations need to be adapted as in [22] where information can be extracted by convolving a 3D kernel to input information. This allows the deep neural network to obtain spatiotemporal features of the input that are useful for complex tasks such as real-time object recognition [36], human action recognition [22] or point cloud labeling [20].

This convolutional operation is typically followed by pooling layers. These take the convolutions' output as input and reduce its dimensionality by aggregation, condensing the representation, and lowering the computational cost [13].

### 3.3 Lifelong learning

Lifelong learning, also referred to as continual learning or incremental learning, is a branch of artificial intelligence that deals with the dynamic nature of data, trying to create models that can adapt to changing environments or scale their capabilities with upcoming knowledge [21].

While there are different types of parameters to consider, there are two types of incremental learning, as follows [33].

- **Task incremental.** The model learns to perform new tasks with time. Hence becoming a multitasking model.
- **Class incremental.** The model learns to classify new classes inside a given task. This allows the creation of models that scale and improve over time, detecting data drifts and adapting to them [42].

However, this type of model faces the challenge of catastrophic forgetting [15], a phenomenon where models may overfit to new data, classes, or tasks, leading to a degradation or forgetting of their performance on previously learned classes and tasks.

To address this issue, various techniques have been proposed. One approach is Elastic weight consolidation, which introduces a regularization term in the loss function to penalize changes to important weights during re-training [25]. Another one is Experience replay, which involves storing and replaying old data while training on new tasks or classes revisiting and learning from past experiences [31].

## 4 OVERALL SOLUTION

This section introduces the proposed approach. The proposal overview (Section 4.1), the process to build the authentication proof (Section 4.2) and the CA model used (Section 4.3) are introduced in the following.

### 4.1 Approach overview

This proposal focuses on studying the feasibility of exploiting the dynamic evolution of MP on streaming data to analyze whether the mere changes of the incremental MP can be used to recognize and authenticate individuals. The process consists of four main phases as depicted in Figure 2:

- (1) **Data extraction.** Sensor data from the accelerometer including X, Y and Z axes are collected from an IoT device. We focus on these features as they have extensively been used for CA purposes [3]. Moreover, as opposed to other behavioral biometrics (e.g., keystrokes, touch gestures), it does not require any explicit action from the user. Such a seamless authentication is beneficial for a healthcare scenario as the one presented before.
- (2) **Data pre-processing and feature extraction/generation.** Sensor data is downsampled, and incremental MP is calculated following a streaming data approach; for every sample, the MP of a buffer is recorded and stored.
- (3) **Dataset preparation.** Once pre-processed, input data is gathered in groups of the size of the used MP window size (WS, recall Section 3). Using these WSs, MPs are computed considering time series of different lengths, referred to as Buffer Size (BS). Thus, BS represents the amount of sensor data stored to compute the MP. Lastly, datasets are prepared for each experiment, splitting into balanced training and testing subsets.
- (4) **Training and assessment.** The model is trained, and performance results, including correct/incorrect authentications are computed.

### 4.2 Authentication proof

When a user  $u_i \in \mathcal{U}$  aims to be authenticated by a server  $\mathcal{S}$ , an authentication proof is built. It contains 3D convolutions based on the incremental MP of the three inputs – X, Y, and Z axes. The process is depicted in Figure 3 and described in the following.

The user device keeps the last BS data readings. Using a given window size WS, the incremental MP (i.e.,  $MP_X(BS_i)$  for the X axis) is computed. This process is sequentially repeated during a period,




Figure 2: General overview

each time using the latest  $BS$  sensor readings – thus, one sensor reading is *forgotten* after each step. For the sake of simplicity, the length of this period (that is, the amount of repetitions) is set to  $WS$ , being  $MP_X(BS_{i+WS-1})$  the last one.

Algorithm 1 shows a pseudocode of this operation. Thus, the authentication proof  $p$  can be regarded as tridimensional data frame formed by three bidimensional arrays – one for each dimension within  $x, y$  and  $z$ . It array has  $WS$  rows, each one being the MP of the latest  $BS$  readings. The size of the MP is  $BS - WS + 1$  by definition [18, 61]. This tridimensional data frame can be visualized as a colorful image by simply interpreting each bidimensional array as one RGB color channel. Therefore, appending all the images, it can be represented as a video <sup>6</sup>.

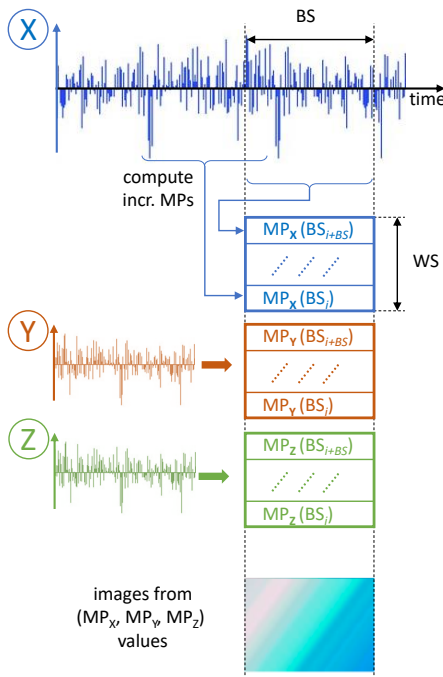


Figure 3: Authentication proof preparation

**Algorithm 1:** Authentication proof  $p$  calculation

**Data:**  $DS_x, DS_y, DS_z$ , accelerometer data of X,Y,Z axes, respectively;  $BS; WS$

**Result:** Authentication proof  $p$

**Initialize:**

$p_x, p_y, p_z \leftarrow \text{null}$  // # To store MPs  
 $p \leftarrow \text{null}$   
 $i \leftarrow 0$

**while**  $i < WS$  **do**

// # We illustrate the x axis here, must be done for y and z too

$MP_i \leftarrow \text{Compute Incremental MP } (MP_x(DS_{x_i}, DS_{x_{i+BS}});$

$p_x \leftarrow \text{append } (MP_i);$

$i \leftarrow i + 1;$

**end**

$p = \text{stack}(p_x, p_y, p_z);$

**4.3 CA model**

The proposed architecture is a classification model (Figure 4) based on a 3D convolution encoder and a Multilayer perceptron (MLP) classifier.

After each convolution, a maxpooling layer is split to reduce the data dimensionality. The resulting features are flattened and sent to a MLP in which the Rectified Linear Activation (ReLU) is applied for being the most used alternative. Moreover, dropout technique helps preventing overfitting [55]. The last layer of this MLP classifier is in charge of defining which class the input belongs to. There are two main variations depending on the scenario (recall Section 2.2). The *single user* scenario uses a sigmoid activation function to specify the class between 0 and 1. In the case of *incremental users* the last layer counts with a softmax activation function in charge of defining multiple classes. In the latter case, the model adopts an incremental approach regarding the number of users. After training in one user, the last layer is expanded with an extra neuron, representing the new added users. Previous model weights are loaded and the model is adapted by replaying the experience of the previous users in addition to the newly added ones.

<sup>6</sup>See two different users: <https://vimeo.com/915165193> and <https://vimeo.com/915166084>

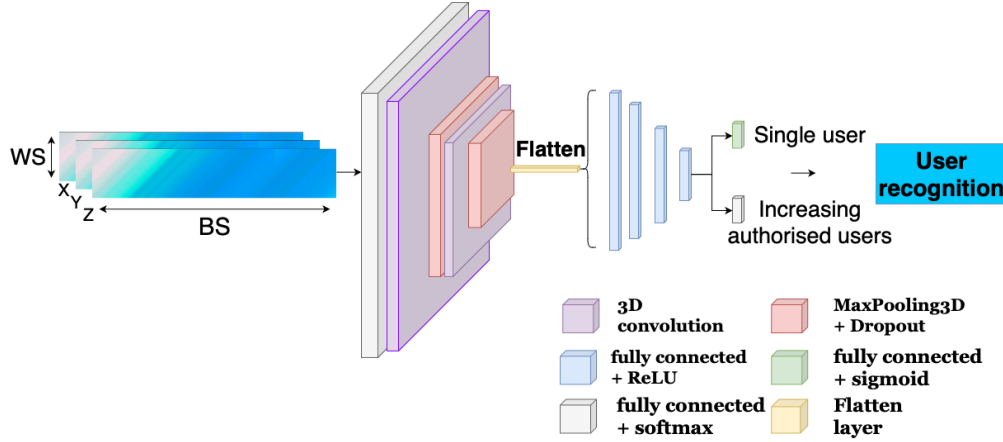


Figure 4: Proposed CA model

Table 1: HMOG dataset sessions

Session ID	Actions
1,7,13,19	Reading + Sitting
2,8,14,20	Reading + Walking
3,9,15,21	Writing + Sitting
4,10,16,22	Writing + Walking
5,11,17,23	Map + Sitting
6,12,18,24	Map + Walking

## 5 EXPERIMENT DESIGN

This Section describes the preparation of the experiments. In particular, Section 5.1 presents the used dataset, Section 5.2 describes the transformations and data preparation and Section 5.3 introduces the experimental settings.

### 5.1 Dataset

HMOG dataset [49] is used in the experiments for being significantly applied for CA purposes (e.g. [52], [14]). This dataset collects multimodal information from smartphone sensors of 100 users during 24 sessions between 5 to 15 minutes. Table 1 summarizes all the IDs as well as the action and situation of the user when recording the sessions.

Following data-saving techniques, only information from the accelerometer is considered. Original data is collected with a sampling rate of 100 Hz. However, the authors point out that 16Hz is sufficient to capture meaningful patterns for CA [49]. Thus, we adopt this sampling rate.

### 5.2 Data preparation

For data preparation, the STUMPY Python library [28] has been used to compute the incremental MP considering different values for BS and WS (recall Section 4.1).

For each data sample on each of the axis (X,Y,Z), incremental MP is computed considering both WS and BS. WS is required to identify the remaining windows to compute the distance to the current one, whereas BS imposes a limit on the size of the data

stream to be considered at a time. The process is repeated with a stride of 1 sample, thus leading to  $WS$  computations of MP. They are normalized considering the range (i.e., maximum and minimum) of values. It must be noted that we compute MP per axis and not as a combination of the three at the same time, as shown in Figure 4.

### 5.3 Experimental settings

Table 2 shows a summary of the parameters of the different experiments. Experiments have been conducted using a NVIDIA 4090 with an Intel i7-10700KF processor and Pytorch [40] as framework to model the training loop. To foster further research, all models are publicly released<sup>7</sup>.

As an optimization, the classical ADAM optimizer was used [24] with a learning rate of  $10^{-3}$  and making use of early stopping as recommended in other works [43]. Regarding batch size, our preliminary tests show that smaller sizes do not affect the results in terms of accuracy. Thus, the batch size has been set following a maximum capacity policy. After a trial and error process a limit of 10 epochs with early stopping and partial saving of model weights are applied, as well as dropout layers with a value of 0.5.

To cover the wide range of actions registered in the HMOG dataset, sessions from 1 to 6 have been selected (recall Table 1). In order to assess the impact of the size of the input data different BSs and WSs are considered. Although BS and WS are measured in terms of data samples (e.g., BS=480 points), for the sake of clarity we use their equivalent in seconds (e.g., BS= 30 s.). In what comes to the split of the dataset, a training/testing ratio of 80/20 has been used for the experiments, using the testing part for validation as well. This leads to an average of 40.56 minutes of training and 10.2 minutes of testing per user.

Beyond the general parameters, there are some particular choices for each scenario. In the *single user* one, the model has been trained using a binary cross-entropy loss function with two classes – the positive class for a random  $u_i$ , the negative one for  $\mathcal{A}$ . For this case,  $\mathcal{A}$  is characterized by 20 random users, using the same total amount of samples as the authorized one. This process is repeated 20 times to ensure the relevance of the results. On the other hand, in

<sup>7</sup>[https://github.com/Luisibear98/ARES-Matrix\\_profile\\_for\\_continuous\\_authentication](https://github.com/Luisibear98/ARES-Matrix_profile_for_continuous_authentication)



**Table 2: Experimental settings summary**

<b>Dataset</b>	HMOG
<b>Buffer Size (BS)</b>	5, 10, 15 and 30 seconds
<b>Windows size (WS)</b>	1, 5 and 15 seconds
<b>Epochs</b>	10
<b>Sampling rate</b>	16 Hz.
<b>Training/Test ratio</b>	80/20
<b>Batch Sizes</b>	Maximum capacity strategy

the *incremental users* scenario the selected loss function is a sparse categorical cross-entropy loss, where each  $u_i$  is assigned a different class and  $\mathcal{A}$  is a single, negative class. In this case, 1, 5, and 10  $u_i$  are randomly chosen. This limit has been chosen since the vast majority of small and medium-sized enterprises in the European Union count on less than 10 employees<sup>8</sup>. On the other hand,  $\mathcal{A}$  is built as in the previous case. Due to the computational requirements of this scenario, 3 repetitions were executed as a trade-off with the relevance.

## 6 EVALUATION

This Section measures the achievement of the established goals (recall Section 2.3). Before that Section 6.1 introduces the metrics at stake. Then, Section 6.2 analyzes the achievement of accuracy (O2) and user scalability (O3) whereas Section 6.3 covers privacy compliance (O1). Lastly, results are discussed in Section 6.4.

### 6.1 Metrics

To assess the model effectiveness, we adopt the typical metrics for the sake of comparability. They are based on combining the amount of True Positives and Negatives (TPs, TNs) and False Positives and Negatives (FPs, FNs). Each one is introduced below:

- **Accuracy.** Gauges the percentage of correctly classified samples:

$$Accuracy = \frac{TP + TN}{TP + TN + FP + FN} \quad (1)$$

- **False Acceptance Rate (FAR).** Counts the amount of times a user is assigned in the class of another one.

$$FAR(\tau) = \rho(Genuine|Fraudulent) = \frac{FP}{FP + TN} \quad (2)$$

- **False Rejection Rate (FRR).** Counts the amount of times a user is not assigned to its correct class:

$$FRR(\tau) = \rho(Fraudulent|Genuine) = \frac{FN}{FN + TP} \quad (3)$$

- **Recall.** Measure the effectiveness in accurately identifying individuals belonging to a particular category.

$$Recall = \frac{TP}{TP + FN} \quad (4)$$

It must be noted that these metrics consider that all errors are the same. Thus, if any  $u_i, u_j$  in the *incremental users* scenario are confused among them, they will be considered as if  $\mathcal{A}$  were classified as legitimate. In some settings both types of errors could have

<sup>8</sup><https://www.statista.com/statistics/878412/number-of-smes-in-europe-by-size/>, last access January 2024.

different severity levels – indeed, confusing authorized users could not be relevant. Therefore, these metrics must be considered as the worst case scenario.

### 6.2 Accuracy and user scalability

Results are depicted in Tables 3 and 4 for the *single user* and *incremental users* scenarios, respectively. Note that results for one  $u_i$  are slightly different in both tables. This is because the underlying AI models are different (recall Section 4.3).

In the following, both goals are assessed. In what comes to accuracy, it requires considering all metrics, as accuracy by itself may lead to unclear results, e.g. missing information about FP and FN. With respect to the user scalability goal, it is only feasible in the *incremental users* scenario.

**6.2.1 Single authorized user  $u_1$ .** In a nutshell, high accuracy and recall are achieved, coupled with low FAR and FRR, until a  $BS = 10$  s. and  $WS = 1$  s.. At this point, the FRR drops to 16%, and the recall for  $u_1$  decreases to 86%.

The trend suggests that increased input information leads to improved results. In the case of a  $BS = 30$  s. and  $WS = 15$  s., the system correctly identifies the authenticated user with 99% accuracy, capturing unique patterns.

With equal BS settings, WS has an impact. For a  $BS = 15$  s., varying WS ( $WS = 15$  s.,  $WS = 5$  s.) yields slightly worse results with a 3-unit decrease in WS. Notably, a  $BS = 10$  s. and  $WS = 5$  s. or  $WS = 1$  s. experiences a more abrupt performance decline, pointing to a balance between buffer information and window size.

FAR holds until  $BS = 5$  s. and  $WS = 1$  s. where it reaches 10%. Correlating with FRR it can be observed that the less information the highest the error. Comparing both values, the system tends to predict  $u_1$  as  $\mathcal{A}$  more frequently than the inverse.

In summary, the more information, the better the results. If less information is given, the model tends to reject  $u_1$  instead of accepting attackers as authorized which is a positive issue in terms of security.

**6.2.2 Incremental authorized users.** The system holds high accuracy and recall until  $BS = WS = 5$  s.. After a certain point, accuracy sharply decreases with an increase in  $\#u_i$ . For instance, when classifying 10 users, the FRR increases to nearly 16%, while the FAR remains low at 1%. This suggests that, as in the previous scenario, the system tends to reject genuine users  $u_i$  but effectively minimizes the acceptance of  $\mathcal{A}$ .

The effect of reducing data depends on the amount of  $u_i$ , as shown in Figure 5. As it can be seen, lower values for BS and WS lead to worse values when the number of  $u_i$  increases.

In terms of users scalability, accuracy holds until  $BS = 15$  s. and  $WS = 5$  s. where recall also begins to be affected. This means that though the system seems to be accurate, it is not good when recognizing new  $u_i$ . In contrast, up to  $BS = WS = 15$  s., the model experiences no decrease, successfully identifying all new  $u_i$ .

**6.2.3 Comparing scenarios.** There some subtle differences between both scenarios. Just results for  $BS = 5$  s. and  $WS = 1$  s. are quite worse in the *incremental users* scenario when  $\#u_i = 5, 10$  – *accuracy* = 67 and *accuracy* = 60, respectively. However, no significant differences are identified in the remaining cases, for




**Table 3: Single user scenario. Results**

<b>BS (s.)</b>	30	15	15	10	5	10	5
<b>WS (s.)</b>	15	15	5	5	5	1	1
<b>Accuracy</b>	<b>99</b>	<b>98</b>	<b>98</b>	<b>97</b>	<b>96</b>	<b>91</b>	<b>81</b>
<b>FAR</b>	0	0	1	0	0	1	10
<b>FRR</b>	1	3	2	5	7	16	26
<b>Recall</b>	<b>99</b>	<b>97</b>	<b>96</b>	<b>94</b>	<b>93</b>	<b>86</b>	<b>78</b>

**Table 4: Incremental users scenario. Results**

<b>#u<sub>i</sub></b>	1	5	10	1	5	10	1	5	10	1	5	10	1	5	10	1	5	10			
<b>BS (s.)</b>	30			15			15			10			5			10			5		
<b>WS (s.)</b>	15			15			5			5			5			1			1		
<b>Accuracy</b>	<b>99</b>	<b>99</b>	<b>99</b>	<b>99</b>	<b>99</b>	<b>99</b>	<b>97</b>	<b>97</b>	<b>97</b>	<b>96</b>	<b>95</b>	<b>94</b>	<b>96</b>	<b>95</b>	<b>89</b>	<b>91</b>	<b>82</b>	<b>72</b>	<b>83</b>	<b>67</b>	<b>60</b>
<b>FAR</b>	0	0	0	0	0	0	0	0	1	2	1	0.5	2	2	1	6	5	4	15	10	6
<b>FRR</b>	0	0	0	0	0	0	1	1	2	2	5	7	2	2	16	11	19	43	15	44	67
<b>Recall</b>	<b>99</b>	<b>99</b>	<b>99</b>	<b>99</b>	<b>99</b>	<b>99</b>	<b>98</b>	<b>98</b>	<b>96</b>	<b>96</b>	<b>94</b>	<b>90</b>	<b>96</b>	<b>96</b>	<b>83</b>	<b>91</b>	<b>79</b>	<b>55</b>	<b>83</b>	<b>54</b>	<b>31</b>



**Figure 5: Incremental users scenario. Accuracy vs #u<sub>i</sub>**

instance, accuracy for  $BS = WS = 15$  s. is beyond 98% in both scenarios.

### 6.3 Privacy compliance

This section assesses the linked minimization requirements imposed to address the privacy goal. Thus, data minimization and storage limitation issues are described separately in the following.

**6.3.1 Data minimization.** The use of MP involves, by itself, the aggregation of users' data to get a "profile" which, in this case, is used for authentication purposes. Raw data is concealed due to the aggregation procedure. Thus, this requirement is met as a direct consequence of using MP [12].

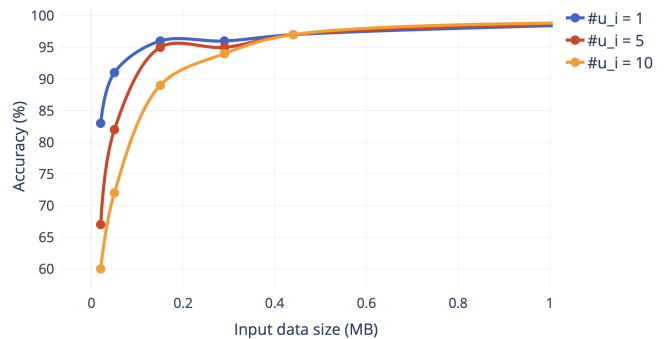
It must be noted that the incremental variant of MP is at stake in this paper – the actual values of MP vary over time as they depend on the evolution of the time series. Therefore, data used in the CA process is ephemeral (with a lifetime of seconds), which contributes to this goal.

**6.3.2 Storage limitation.** The amount of information at stake is depicted in Table 5, including the model size and the input data size. The former is relevant for the CA server, which has to store

the model to make authentication decisions. The latter is important from the perspective of the user to be authenticated – sacrificing storage space for a security service does not seem to be a natural priority.

In what comes to the model size, it is virtually the same in both scenarios. Although the models are different, the storage size is pretty similar with negligible variations. The largest memory footprint of this approach is due to the size of the neural network model which on its highest is about 1.78 GB. As expected, there is a positive correlation among model size and input data size – the model needs to learn more information when more input data is provided. In what comes to the input data, the configuration with the most significant memory demand is  $BS = 30$  s., capping at a maximum of 2.64 MB. At the light of these figures, the approach is suitable for both server and end-user devices considering their typical storage capabilities.

As a means to find the optimal storage size while achieving satisfactory accuracy and user scalability, Figure 6 visually depicts the evolution of accuracy and recall relative to the input data size. For the sake of readability, input data size is limited to 1 MB as greater values lead to negligible performance improvements.



**Figure 6: Accuracy and user scalability vs input data storage**

**Table 5: Storage requirements**

<b>BS (s.)</b>	30	15	15	10	5	10	5
<b>WS (s.)</b>	15	15	5	5	5	1	1
<b>Model Size</b>	1.78 GB	917.85 MB	307.85 MB	207.85 MB	160.85 MB	39.85 MB	19.85 MB
<b>Auth. proof (<math>p</math>) size</b>	<b>2.64 MB</b>	<b>1.32 MB</b>	<b>0.44 MB</b>	<b>0.29 MB</b>	<b>0.15 Mb</b>	<b>0.05 MB</b>	<b>0.02 MB</b>

Within the range of 0 MB to 0.3 MB, there is a pronounced increase in both accuracy and recall performance. However, as the information size surpasses this range, the rate of improvement becomes less substantial. A convergence is observed around 0.4 MB, which suggests that the optimal configuration for both scenarios is  $BS = 15$  s. and  $WS = 5$  s.

## 6.4 Discussion

Our results confirm that the use of MP is a suitable technique for CA purposes. Interestingly, it not only addresses several GDPR-based concerns, but also achieves promising performance ratios under some settings.

Despite the overall analysis, there are three experimental issues that deserve our attention. On the one hand, the training-testing split rate has been fixed according to widespread practices. In practice, they require some time for the onboarding process. While it is affordable, different splits (e.g., 60-40 or even 20-80) could offer interesting results. Our preliminary results show that the training period could be shortened, but it requires a comprehensive assessment process that is left to future work. The second issue is related to selecting  $WS$  as the parameter to set the amount of convolutions. Even if it leads to very limited storage needs, using lower values could contribute to the suitability of this mechanism to resource-constrained devices (e.g., wearables or implantable devices). Lastly, the length of the authentication tests (around 10 minutes per user, recall Section 5.3) is suitable for achieving *continuous* authentication, the long-term suitability of this approach remains unexplored. However, to the best of authors' knowledge, there are no longer datasets with high-resolution accelerometer data.

On the other hand, our results suggest that this approach cannot be easily integrated into some environments that may have specific movement patterns. For example, industrial environments may have additional constraints in what comes to movement. Similarly, moving surfaces like trains may also pose challenges to the adoption of this technique.

## 7 RELATED WORK

CA is a field of active research in the last years. Many CA approaches have leveraged raw data from the sound [9], touch dynamics [46, 47, 50, 57] or sensors [2, 8, 19, 49] to capture meaningful features that allow identifying individuals. However, despite these efforts, privacy has yet to be significantly considered, worsening the impact of data leakage on a system that could be based on this type of feature to authenticate.

Chauhan *et al.* [9] and Shen *et al.* [47] work in incremental learning approaches that could be efficient for online scenarios. The former leverages breath data and the latter touch dynamics. Nonetheless, they do not contemplate any type of privacy preservation beyond on-device execution of the classification systems.

Indeed, there is a line of research to enhance privacy while leveraging continuous authentication based on behavior [16]. One of the approaches is federated learning, where Wazzeah *et al.* [54] proposed a federated scenario where each of the clients shares a portion of data to a central server, which creates a global warm-up model that is later distributed with each of the clients. Privacy via this approach is based on how data is shared more than how the data is processed, in contrast to our approach. For instance, this architecture could be used alongside our proposal to enhance privacy. By contrast, Sanchez *et al.* [45] reach privacy via aggregation, filtering, and transformation of the features. They proposed a multi-device cloud-based system in which users are authenticated based on interaction among devices.

Other lines of research work on studying how to encrypt the data efficiently, such as data being transferred safely between the authentication servers and the endpoints [7, 19] which perform predictions over encrypted data. However, this methodology requires maintaining a set of users and keys to encrypt the data, and then, each of the classifiers would need to be trained for each of the keys and every time the key changes. In addition, the first proposal applies classic support vector machines (SVM), and the latter, a function that compares the new samples with the existing ones, makes decisions through cosine similarities.

On the other hand, previous efforts have not applied MP for continuous authentication, although MP has already been explored as an additional feature for anomaly detection on images [32], as a sole feature to classify mammals [38] or for detecting denial of service attacks [5]. Some classification efforts consider MP in its online streaming data variants. [62] detects seismic motifs in online streaming data with a GPU-efficient algorithm, but they do not explore the integration with deep learning to enhance or add an extra layer to classify the detected motifs. Another line of research is focused on analyzing and taking advantage of the extra layer of privacy that can be obtained from MP by using it as a way to guide algorithms to generate synthetic time series that preserve specific MPs [12] or to improve MP to create more resilient algorithms [59] to minimize the leakage of information that can be inferred from the MPs indexes. Therefore, the previous efforts have not applied MP for continuous authentication.

Table 6 shows an overview of the related works, where any of them addresses our same goals. However, it is remarkable that [19] and [7] apply encryption and they are highlighted as partially private because they do not provide privacy in the sense of data and storage minimization and [45] does not deal with storage minimization neither.

**Table 6: Related work.** ✓=Addressed; P=Partially; ×=Not addressed

	<b>Sitová et al.</b> [49]	<b>Centeno et al.</b> [8]	<b>Chauhan et al.</b> [9]	<b>Sánchez et al.</b> [45]	<b>Hernández-Álvarez et al.</b> [19]	<b>Baig et al.</b> [7]	<b>Stylios et al.</b> [50]	<b>Shen et al.</b> [47]	<b>Ours</b>
<b>Year</b>	2015	2017	2020	2020	2021	2023	2023	2023	2024
<b>Technique</b>	Scaled Manhattan (SM), Scaled Euclidian (SE), SVM.	Autoencoders	DCN	MLP, XGBoost, RF and LSTM	SVM	Similarity based model	LSTM	LSTM	DCN
<b>Dataset</b>	HMOG	Lab-environment dataset / crowdsignals.io	breathing dataset	Own dataset	Sherlock database	Free vs. transcribed and scrolling interactions datasets	Own dataset	Own dataset	HMOG
<b>Features</b>	HMOG, keystroke, and tap features	Accelerometer	Sound data	Multiple sensors from multiple-devices	Accelerometer and gyroscope	Swipe gesture and key-stroke dynamics	Fusion of Keystrokes dynamics and touch gestures	Touch dynamics data	Accelerometer
<b>Metrics</b>	EER	EER	Accuracy	Accuracy	Accuracy, EER	EER	Accuracy, EER	Accuracy, EER	Accuracy
<b>Results</b>	13.62%	2.2%	97%	99.32%	76.84 (Accuracy)% / 23.24% (EER)	12% / 20%	99% (Accuracy) / 1% (EER)	95% 99% (Accuracy) / 5% (EER)	99%
<b>Privacy compliance (O1)</b>	×	×	×	P	P	P	×	×	✓
<b>Data minimization (O1.1)</b>	×	×	×	✓	×	×	×	×	✓
<b>Storage Minimization (O1.2)</b>	×	×	×	×	×	×	×	×	✓
<b>Accuracy (O2)</b>	×	✓	✓	✓	×	×	✓	✓	✓
<b>User scalability (O3)</b>	×	×	✓	×	×	×	×	✓	✓

## 8 CONCLUSION

This paper has proposed a Continuous Authentication (CA) model that copes with two opposing requirements – being adaptive to the amount of authorized users over time while minimizing the data at stake so as to protect the user privacy. The approach has leveraged incremental Matrix Profile and Deep Learning applied to accelerometer data. Results show that the model offers promising performance figures both in the single and multi-user settings, while imposing affordable storage requirements.

Our results open up a number of future research directions. On the one hand, characterizing the maximum amount of users that can be authorized may be relevant for scenarios with high worker volatility. On the other hand, machine unlearning techniques could be applied for those scenarios in which workers are substituted, thus no longer authorized. Lastly, the analysis of the suitability of this approach to other behavioral traits is useful to assess the widespread application of our approach. An implicit assumption made in our study is that MP helps enforcing privacy [12, 59]. While our focus here was on data minimisation and storage limitation, an important future work is to precisely quantify the privacy provided by MP.

## ACKNOWLEDGMENTS

The work of Luis Ibanez-Lissen has been supported by INCIBE grant APAMciber within the framework of the Recovery, Transformation and Resilience Plan funds, financed by the European Union (Next Generation). The work of Jose Maria de Fuentes is supported by the Madrid Government (Comunidad de Madrid-Spain) under the Multiannual Agreement with UC3M ('Fostering Young Doctors Research', DEPROFAKE-CM-UC3M), and in the context of the V PRICIT (Research and Technological Innovation Regional Programme). Additionally, Lorena Gonzalez and Jose M. de Fuentes have received support from UC3M's Requalification programme, funded by the Spanish Ministerio de Universidades with EU recovery funds. Nicolas Anciaux is supported by a grant from ANR, the French national research agency, as part of the PEPR Cybersecurité, Interdisciplinary Project on Privacy iPOP (ANR-22-PECY-0002).

## CONFLICT OF INTEREST

Authors declare they have no conflict of interest.

## DATA AVAILABILITY STATEMENT

Data is publicly available at: [https://github.com/Luisibear98/ARES-Matrix\\_profile\\_for\\_continuous\\_authentication](https://github.com/Luisibear98/ARES-Matrix_profile_for_continuous_authentication)

## REFERENCES

- [1] Kumar Abhishek, Sahana Roshan, Prabhat Kumar, and Rajeev Ranjan. 2013. A comprehensive study on multifactor authentication schemes. In *Advances in Computing and Information Technology: Proc. Second International Conference on Advances in Computing and Information Technology (ACITY) July 13-15, 2012, Chennai, India-Volume 2*. Springer, 561–568.
- [2] Mohammed Abuhamad, Tamer Abuhmed, David Mohaisen, and DaeHun Nyang. 2020. AUToSen: Deep-learning-based implicit continuous authentication using smartphone sensors. *IEEE Internet of Things Journal* 7, 6 (2020), 5008–5020.
- [3] Mohammed Abuhamad, Ahmed Abusnaina, DaeHun Nyang, and David Mohaisen. 2020. Sensor-based continuous authentication of smartphones' users using behavioral biometrics: A contemporary survey. *IEEE Internet of Things Journal* 8, 1 (2020), 65–84.
- [4] Fatimah Hussain Al-Naji and Rachid Zagrouba. 2020. A survey on continuous authentication methods in Internet of Things environment. *Computer Communications* 163 (2020), 109–133.
- [5] Faisal Alotaibi and Alexei Lisitsa. 2021. Matrix profile for DDoS attacks detection. In *2021 16th Conference on Computer Science and Intelligence Systems (FedCSIS)*. IEEE, 357–361.
- [6] Nicolas Anciaux, Philippe Bonnet, Luc Bouganim, Benjamin Nguyen, Philippe Pucheral, Iulian Sandu Popa, and Guillaume Scerri. 2019. Personal Data Management Systems: The security and functionality standpoint. *Inf. Syst.* 80 (2019), 13–35. <https://doi.org/10.1016/j.is.2018.09.002>
- [7] Ahmed Fraz Baig, Sigurd Eskeland, and Bian Yang. 2023. Privacy-preserving continuous authentication using behavioral biometrics. *International Journal of Information Security* 22, 6 (2023), 1833–1847.
- [8] Mario Parreño Centeno, Aad van Moorsel, and Stefano Castruccio. 2017. Smartphone continuous authentication using deep learning autoencoders. In *2017 15th annual conference on privacy, security and trust (pst)*. IEEE, 147–1478.
- [9] Jagmohan Chauhan, Young D Kwon, Pan Hui, and Cecilia Mascolo. 2020. Con-tauth: Continual learning framework for behavioral-based user authentication. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 4, 4 (2020), 1–23.
- [10] Long Cheng, Fang Liu, and Danfeng Yao. 2017. Enterprise data breach: causes, challenges, prevention, and future directions. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery* 7, 5 (2017), e1211.
- [11] Dieter De Paepe, Diego Nieves Avendano, and Sofie Van Hoecke. 2019. Implications of z-normalization in the matrix profile. In *International Conference on Pattern Recognition Applications and Methods*. Springer, 95–118.
- [12] Audrey Der, Chin-Chia Michael Yeh, Yan Zheng, Junpeng Wang, Huiyuan Chen, Zhongfang Zhuang, Liang Wang, Wei Zhang, and Eamonn Keogh. 2023. Time series synthesis using the matrix profile for anonymization. In *2023 IEEE International Conference on Big Data (BigData)*. IEEE, 1908–1911.
- [13] Hossein Gholamalinezhad and Hossein Khosravi. 2020. Pooling methods in deep neural networks, a review. *arXiv preprint arXiv:2009.07485* (2020).
- [14] Giacomo Giorgi, Andrea Saracino, and Fabio Martinelli. 2021. Using recurrent neural networks for continuous authentication through gait analysis. *Pattern Recognition Letters* 147 (2021), 157–163.
- [15] Ian J Goodfellow, Mehdi Mirza, Da Xiao, Aaron Courville, and Yoshua Bengio. 2013. An empirical investigation of catastrophic forgetting in gradient-based neural networks. *arXiv preprint arXiv:1312.6211* (2013).
- [16] Kimmo Halunen and Visa Vallivaara. 2016. Secure, usable and privacy-friendly user authentication from keystroke dynamics. In *Secure IT Systems: 21st Nordic Conference, NordSec 2016, Oulu, Finland, November 2-4, 2016. Proceedings 21*. Springer, 256–268.
- [17] Zecheng He, Tianwei Zhang, and Ruby B Lee. 2019. Model inversion attacks against collaborative inference. In *Proceedings of the 35th Annual Computer Security Applications Conference*. 148–162.
- [18] Colin Hehir and Alan F Smeaton. 2023. Calculating the matrix profile from noisy data. *Plos one* 18, 6 (2023), e0286763.
- [19] Luis Hernández-Álvarez, José María De Fuentes, Lorena González-Manzano, and Luis Hernández Encinas. 2021. SmartCAMPP-Smartphone-based continuous authentication leveraging motion sensors with privacy preservation. *Pattern Recognition Letters* 147 (2021), 189–196.
- [20] Jing Huang and Suya You. 2016. Point cloud labeling using 3d convolutional neural network. In *2016 23rd International Conference on Pattern Recognition (ICPR)*. IEEE, 2670–2675.
- [21] Chip Huyen. 2022. *Designing machine learning systems*. " O'Reilly Media, Inc."
- [22] Shuiwang Ji, Wei Xu, Ming Yang, and Kai Yu. 2012. 3D convolutional neural networks for human action recognition. *IEEE transactions on pattern analysis and machine intelligence* 35, 1 (2012), 221–231.
- [23] Eamonn Keogh and Shruti Kasetty. 2002. On the need for time series data mining benchmarks: a survey and empirical demonstration. In *Proceedings of the eighth ACM SIGKDD international conference on Knowledge discovery and data mining*. 102–111.
- [24] Diederik P. Kingma and Jimmy Ba. 2014. Adam: A Method for Stochastic Optimization.
- [25] James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, et al. 2017. Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences* 114, 13 (2017), 3521–3526.
- [26] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. 2012. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems* 25 (2012).
- [27] Riad Ladjel, Nicolas Anciaux, Philippe Pucheral, and Guillaume Scerri. 2019. A Manifest-Based Framework for Organizing the Management of Personal Data at the Edge of the Network. In *Information Systems Development: Information Systems Beyond 2020, ISD 2019 Proceedings, Toulon, France, August 28-30, 2019*, Alena Sitarheeva, Chris Barry, Michael Lang, Henry Linger, and Christoph Schneider (Eds.). ISEN Yncréa Méditerranée / Association for Information Systems. <https://aisel.aisnet.org/isd2014/proceedings2019/ManagingISD/1>
- [28] Sean M. Law. 2019. STUMPY: A Powerful and Scalable Python Library for Time Series Data Mining. *The Journal of Open Source Software* 4, 39 (2019), 1504.
- [29] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. 2015. Deep learning. *nature* 521, 7553 (2015), 436–444.
- [30] Yann LeCun, Bernhard Boser, John S Denker, Donnie Henderson, Richard E Howard, Wayne Hubbard, and Lawrence D Jackel. 1989. Backpropagation applied to handwritten zip code recognition. *Neural computation* 1, 4 (1989), 541–551.
- [31] Long-Ji Lin. 1992. Self-improving reactive agents based on reinforcement learning, planning and teaching. *Machine learning* 8 (1992), 293–321.
- [32] Qian Liu, Carson K Leung, and Pingzhao Hu. 2020. A two-dimensional sparse matrix profile DenseNet for COVID-19 diagnosis using chest CT images. *IEEE Access* 8 (2020), 213718–213728.
- [33] Vincenzo Lomonaco, Lorenzo Pellegrini, Andrea Cossu, Antonio Carta, Gabriele Graffieti, Tyler L Hayes, Matthias De Lange, Marc Masana, Jary Pomponi, Gido M Van de Ven, et al. 2021. Avalanche: an end-to-end library for continual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 3600–3610.
- [34] Kalle Lyytinen and Youngjin Yoo. 2002. Ubiquitous computing. *Commun. ACM* 45, 12 (2002), 63–96.
- [35] Frank Madrid, Shima Imani, Ryan Mercer, Zachary Zimmerman, Nader Shakibay, and Eamonn Keogh. 2019. Matrix profile xx: Finding and visualizing time series motifs of all lengths using the matrix profile. In *2019 IEEE International Conference on Big Knowledge (ICBK)*. IEEE, 175–182.
- [36] Daniel Maturana and Sebastian Scherer. 2015. Voxnet: A 3d convolutional neural network for real-time object recognition. In *2015 IEEE/RSJ international conference on intelligent robots and systems (IROS)*. IEEE, 922–928.
- [37] Fabian Monrose and Aviel Rubin. 1997. Authentication via keystroke dynamics. In *Proceedings of the 4th ACM Conference on Computer and Communications Security*. 48–56.
- [38] Erin C Murnane and NAVAL RESEARCH LAB WASHINGTON DC. 2021. Detection, Classification, and Prediction of Satellite Tagged Marine Mammals for the US Navy. (2021).
- [39] Isao Nakanishi, Sadanao Baba, and Chisei Miyamoto. 2009. EEG based biometric authentication using new spectral features. In *2009 International Symposium on Intelligent Signal Processing and Communication Systems (ISPACS)*. IEEE, 651–654.
- [40] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. 2019. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems* 32 (2019).
- [41] David A Petti. 1998. An Argument for the Implementation of a Biometric Authentication System (BAS). *J. Pat. & Trademark Off. Soc'y* 80 (1998), 703.
- [42] Sylvestre-Alvise Rebuffi, Alexander Kolesnikov, Georg Sperl, and Christoph H Lampert. 2017. icarl: Incremental classifier and representation learning. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*. 2001–2010.
- [43] Sebastian Ruder. 2016. An overview of gradient descent optimization algorithms.
- [44] Riseul Ryu, Soonja Yeom, Soo-Hyung Kim, and David Herbert. 2021. Continuous multimodal biometric authentication schemes: a systematic review. *IEEE Access* 9 (2021), 34541–34557.
- [45] Pedro Miguel Sánchez Sánchez, Lorenzo Fernández Maimó, Alberto Huertas Celdrán, and Gregorio Martínez Pérez. 2021. AuthCODE: A privacy-preserving and multi-device continuous authentication architecture based on machine and deep learning. *Computers & Security* 103 (2021), 102168.
- [46] Muhammad Shahzad, Alex X Liu, and Arjmand Samuel. 2016. Behavior based human authentication on touch screen devices using gestures and signatures. *IEEE Transactions on Mobile Computing* 16, 10 (2016), 2726–2741.
- [47] Zhihao Shen, Shun Li, Xi Zhao, and Jianhua Zou. 2023. InCreAuth: Incremental Learning based Behavioral Biometric Authentication on Smartphones. *IEEE Internet of Things Journal* (2023).
- [48] Karen Simonyan and Andrew Zisserman. 2014. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556* (2014).

- [49] Zdeňka Sitová, Jaroslav Šeděnka, Qing Yang, Ge Peng, Gang Zhou, Paolo Gasti, and Kiran S Balagani. 2015. HMOG: New behavioral biometric features for continuous authentication of smartphone users. *IEEE Transactions on Information Forensics and Security* 11, 5 (2015), 877–892.
- [50] Ioannis Stylios, Sotirios Chatzis, Olga Thanou, and Spyros Kokolakis. 2023. Continuous Authentication with Feature-Level Fusion of Touch Gestures and Key-stroke Dynamics to Solve Security and Usability Issues. *Computers & Security* (2023), 103363.
- [51] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. 2015. Going deeper with convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 1–9.
- [52] Hasan Can Volaka, Gulfem Alptekin, Okan Engin Basar, Mustafa Isbilen, and Ozlem Durmaz Incel. 2019. Towards continuous authentication on mobile phones using deep learning models. *Procedia Computer Science* 155 (2019), 177–184.
- [53] Nikola Vukobrat, Nemanja Maček, Saša Adamović, Muzafer Saračević, and Milan Gnjatović. 2023. Implementation of two factor authentication using face and iris biometrics. In *Blockchain Technology Solutions for the Security of IoT-Based Healthcare Systems*. Elsevier, 77–96.
- [54] Mohamad Wazzeh, Hakima Ould-Slimane, Chamseddine Talhi, Azzam Mourad, and Mohsen Guizani. 2022. Privacy-preserving continuous authentication for mobile and iot systems using warmup-based federated learning. *IEEE Network* (2022).
- [55] Haibing Wu and Xiaodong Gu. 2015. Towards dropout training for convolutional neural networks. *Neural Networks* 71 (2015), 1–10.
- [56] Chin-Chia Michael Yeh, Yan Zhu, Liudmila Ulanova, Nurjahan Begum, Yifei Ding, Hoang Anh Dau, Diego Furtado Silva, Abdullah Mueen, and Eamonn Keogh. 2016. Matrix profile I: all pairs similarity joins for time series: a unifying view that includes motifs, discords and shapelets. In *2016 IEEE 16th international conference on data mining (ICDM)*. Ieee, 1317–1322.
- [57] Ahmad Zairi Zaidi, Chun Yong Chong, Zhe Jin, Rajendran Parthiban, and Ali Safaa Sadiq. 2021. Touch-based continuous mobile device authentication: State-of-the-art, challenges and opportunities. *Journal of Network and Computer Applications* 191 (2021), 103162.
- [58] Matthew D Zeiler and Rob Fergus. 2013. Visualizing and Understanding Convolutional Networks. arXiv:1311.2901 [cs.CV]
- [59] Li Zhang, Jiahao Ding, Yifeng Gao, and Jessica Lin. 2023. PMP: Privacy-Aware Matrix Profile against Sensitive Pattern Inference for Time Series. In *Proceedings of the 2023 SIAM International Conference on Data Mining (SDM)*. SIAM, 891–899.
- [60] Yan Zhu, Shaghayegh Gharghabi, Diego Furtado Silva, Hoang Anh Dau, Chin-Chia Michael Yeh, Nader Shakibay Senobari, Abdulaziz Almaslukh, Kaveh Kamgar, Zachary Zimmerman, Gareth Funning, et al. 2020. The Swiss army knife of time series data mining: ten useful things you can do with the matrix profile and ten lines of code. *Data Mining and Knowledge Discovery* 34 (2020), 949–979.
- [61] Yan Zhu, Chin-Chia Michael Yeh, Zachary Zimmerman, Kaveh Kamgar, and Eamonn Keogh. 2018. Matrix profile XI: SCRIMP++: time series motif discovery at interactive speeds. In *2018 IEEE International Conference on Data Mining (ICDM)*. IEEE, 837–846.
- [62] Yan Zhu, Zachary Zimmerman, Nader Shakibay Senobari, Chin-Chia Michael Yeh, Gareth Funning, Abdullah Mueen, Philip Brisk, and Eamonn Keogh. 2016. Matrix profile ii: Exploiting a novel algorithm and gpus to break the one hundred million barrier for time series motifs and joins. In *2016 IEEE 16th international conference on data mining (ICDM)*. IEEE, 739–748.
- [63] Yan Zhu, Zachary Zimmerman, Nader Shakibay Senobari, Chin-Chia Michael Yeh, Gareth Funning, Abdullah Mueen, Philip Brisk, and Eamonn Keogh. 2018. Exploiting a novel algorithm and GPUs to break the ten quadrillion pairwise comparisons barrier for time series motifs and joins. *Knowledge and Information Systems* 54 (2018), 203–236.
- [64] Jingwei Zuo, Karine Zeitouni, and Yehia Taher. 2019. Incremental and adaptive feature exploration over time series stream. In *2019 IEEE international conference on big data (Big Data)*. IEEE, 593–602.